 PROCEEDINGS OF THE IEEE, VOL. 66, NO. 11, NOVEMBER 1978

Applications of Information Networks

J. C. R. LICKLIDER MEMBER, IEEE, AND ALBERT VEZZA, MEMBER, IEEE

Invited Paper

Abstract—The present and projected applications of computer-communication networks or information networks include electronic mail, teleconferencing, "the office of the future," management information systems, modeling, "computerized commerce," monitoring of patients, military command and control, home security, education, and news. This paper briefly examines 30 such applications and the network capabilities they require. It presents a way of estimating the relative importance of various network characteristics and of predicting the suitability of a network or network architecture for a given set of applications. The paper then considers several issues that relate to the political, social, and economic impacts of networks. Among the issues are privacy, security, compatibility, impact on productivity, the roles of networks in international technology transfer and economic competition, and the confluence or collision of the fields of computers and telecommunications.

I. Introduction

THE SUBJECT of this paper is applications of networks. The networks involve the use of computers, but computation in the narrow sense does not necessarily dominate the applications. The scope of the paper includes, no less than computation, computer-based applications in which the main emphasis is on communication among people, on access to information, or on control of systems, organizations, or—to mention early one of the deepest though least imminent concerns—societies. The applications of networks that we shall examine include electronic message communication [1]–[4], electronic funds transfer [5], access to information, computer-based office work and "telework," management of organizations and command and control of operations, education, entertainment and recreation, reservations and ticketing, and several others.

Some of the problems and issues in network applications are mainly technical and some are mainly nontechnical, but almost all are mixtures of the two, and in most of them the technical and nontechnical factors interact strongly. For example, the relative merits of circuit switching and packet switching are mainly a technical matter, but the fact that the electronic switching stations of the existing telecommunications "plant" are circuit switches surely is an economic factor in the circuit-switching/packet-switching issue. The determination of what should be the individual citizen's right to informational privacy is mainly a nontechnical matter, but the pragmatics of providing informational security, the technical basis for assurance of privacy, must enter the decision process. The national telecommunications plant-in-being of the U.S. figures strongly in many of these problems and issues and forces them to involve both technical and nontechnical factors. The plant is valued at something like $120 billion, and most of it was designed to carry analog voice signals, which are quite different in their spectral and temporal parameters and in their requirements for error handling and security, from digital computer signals of the kinds that will flow through the networks of the future. Because of its inherent redundancy, speech remains intelligible even when mixed with considerable amounts of noise, but even a single undetected error—a single bit—can have extremely serious consequences in electronic funds transfer (EFT) or seriously degrade the performance of a network carrying enciphered information.

One of the major motivations for networking is the need to share resources. The main resources that are often advantageous to share are communications facilities, computer facilities, and information itself. The design of a network can make it easier or more difficult to share resources and thus directly influence the amount of resource sharing that will occur. The amount of communication facilities sharing depends upon many design factors, all of which influence how well the network is able to allocate resources dynamically in response to changing needs and availabilities. Though the need for sharing certain types of computational facilities may diminish with the arrival of the age of the personal computer, it is not at all likely that the need to share resources will disappear altogether. Geographically distributed users can share, through a computer network, the costly high-performance computers that are required to solve certain large computational problems. Even those using personal computers to satisfy the bulk of their computing needs may wish to avail themselves through a network of special software services provided by vendors—and they will certainly wish to communicate with one another.

The sharing of information is the most important type of resource sharing. The term "information sharing" immediately conjures up the thought of sharing large data banks of information among many users, but that is only one aspect of information sharing. All the applications discussed in this paper have aspects of information sharing. Applications concerned with communications, management, commerce, government, protection, education, and awareness all involve sharing. The convenience and effectiveness with which sharing can be accomplished and the facility with which information can flow across the boundaries of individual application programs will have profound effects on how well the applications serve their intended purposes.

Many problems and issues arise from the interaction of information sharing with information security. For example, should EFT have a network or networks of its own to simplify the problem of providing secure transmission, processing, and storage of funds data, or should EFT messages be carried over a general-purpose network so that a reservations and ticketing operation can be completed in a single transaction involving the traveler's organization, the airline and the bank.

Manuscript received March 10, 1978; revised July 7, 1978.
The authors are with the Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

II. Applications

In the context of information networks, just as in the context of computer systems, an application is essentially the implementation of a purpose. Applications, like purposes, may be defined narrowly or broadly. When the airlines began to develop computer-based reservations systems and formed a consortium, ARINC, to interconnect several of their systems, the application was narrow: airline reservations. Now, after years of growth and augmentation, one can rent a car, reserve a hotel room, and arrange to be greeted with flowers and mariachi music. The broadened application might be defined as reservations and ticketing for almost anything that flies to or can be purchased at a distant place. One can project the broader definition into the future and envision a general reservations and ticketing application, operating in a nationwide or worldwide common-carrier network, through which anyone could examine the availability of, and reserve or buy a ticket to, almost anything in the broad class to which reservations and tickets apply. But, of course, there is no reason to stop at that particular class. One can expand the scope further and arrive at "computerized commerce," dealing with the whole gamut of things that can be bought and sold. The very broadly defined application would include advertising, dynamic pricing, and computer-based purchasing strategies. It might even make a place for cartels of suppliers and cooperatives of consumers. No doubt there would be vigorous competition among several or many offerers of the application, and perhaps one can imagine even a "meta-market," an over-arching system that interconnects and integrates the competing "computerized commerce applications."

In any event, that introduces the notion of applications of information networks. It might serve to introduce, also, the notion of issues, which involve the interplay of the opportunity and the threat aspects of applications. It is not difficult to imagine the mischief that could be played by pranksters or dissidents in a poorly protected, publicly accessible, nationwide reservations system.

A. Basic Applications

Computer-communication networks perform three basic classes of operations upon information: transmission, processing, and storage. The earliest recognized applications of networks were essentially separate exploitations of the three basic classes of operations. Transmission of information through a network from a program in one computer to a program in another, of course, requires some processing and some storage (memory), but in simple message communication and file transfer interest is focused sharply on transmission. In every practical computer, processing requires storage (memory or registers), but in early time-sharing services such as Quiktran [6]–[8], which when introduced did not provide intersession file storage, the (dial telephone) network application was essentially access to processing. In the Datacomputer [9] service available through the ARPANET [10]–[12], although processing is involved in both storage and retrieval, one of the main applications is essentially access to storage in and of itself: the Datacomputer is a place to park bits.

A fourth essential network function combines the basic transmission, processing, and storage operations to provide access to information—with the focus of interest on the information itself, rather than on any of the three basic elementary operations.

Simple message communication and file transfer, access to time-shared processing, access to storage, and access to information are important as well as fundamental network functions, but they are no longer typical of the activities or services we associate with the term "application." In present-day parlance, "application" suggests something more highly differentiated and specialized and closer to some specific task or mission.

B. Communication Applications

In the developmental history of the ARPANET, electronic message service was a sleeper. Even before the network included a dozen computers, several message programs were written as natural extensions of the "mail" systems that had arisen in individual time-sharing systems in the early 1960's. By the Fall of 1973, the great effectiveness and convenience of such fast, informal message services as SNDMSG [13] had been discovered by almost everyone who had worked on the development of the ARPANET—and especially by the then Director of ARPA, S. J. Lukasik, who soon had most of his office directors and program managers communicating with him and with their colleagues and their contractors via the network. Thereafter, both the number of (intercommunicating) electronic mail systems and the number of users of them on the ARPANET increased rapidly.

1) Electronic Mail, Electronic Message Systems: It soon became obvious that the ARPANET was becoming a human-communication medium with very important advantages over normal U.S. mail and over telephone calls. One of the advantages of the message systems over letter mail was that, in an ARPANET message, one could write tersely and type imperfectly, even to an older person in a superior position and even to a person one did not know very well, and the recipient took no offense. The formality and perfection that most people expect in a typed letter did not become associated with network messages, probably because the network was so much faster, so much more like the telephone. Indeed, tolerance for informality and imperfect typing was even more evident when two users of the ARPANET linked their consoles together and typed back and forth to each other in an alphanumeric conversation. Among the advantages of the network message services over the telephone were the fact that one could proceed immediately to the point without having to engage in small talk first, that the message services produced a preservable record, and that the sender and receiver did not have to be available at the same time. A typical electronic mail system now provides a rudimentary editor to facilitate preparation of messages, a multiple-addressee feature to make it easy to send the same message to several people, a file-inclusion scheme to incorporate already prepared text files into a message, an alerting mechanism to tell the user that he has new mail in his mailbox, facilities for reading received messages, and a "help" subsystem. The prospects of electronic mail appear to have caught the attention of computer manufacturers and software and time-sharing firms as well as telephone companies, national telecommunication authorities, and the U.S. Post Office—and most of them now seem to be planning, developing, or even offering some kind of electronic mail service.

Even before electronic mail was well established, it had become apparent that users would need computer aids for scanning, indexing, filing, retrieving, summarizing, and responding to messages. Indeed messages are usually not isolated documents but documents prepared and transmitted in the

course of performing complex activities often called "tasks." Within task contexts, messages are related to other messages and to documents of other kinds, such as forms and reports. It seems likely that we shall see a progressive escalation of the functionality and comprehensiveness of computer systems that deal with messages. If "electronic mail" refers to an early stage in the progression, "electronic message system" is appropriate for a later stage and "computer-based office system" or some comparable term for the stage of full integration. At some intermediate point, message service will no doubt be blended with direct user-to-user linking to provide for delay-free conversation whenever both sender and receiver are on-line at the same time and prefer conversation to sequential exchange of messages.

2) Duologue and Teleconferencing: Although there has not been, thus far, very much use of networks for one-on-one interaction between users, it seems likely that some kind of computer-augmented two-person telephone communication will one day be one of the main modes of networking. In order to displace the conventional telephone, "teleduologue" will probably have to offer speech, writing, drawing, typing, and possibly some approximation to television, all integrated into a synergic pattern with several kinds of computer support and facilitation. The two communicators (and their supporting programs) will then be able to control displays in certain areas of each other's display screens and processes in certain sectors of each other's computers. Throughout a duologue, each communicator will be advised by his own programs and will use information from his own data bases and other sources accessible to him. The effect will be to provide each communicator with a wide choice of media for each component of his communication and with a very fast and competent supporting staff.

A teleconference [14]–[17] is an organized interaction, through a communication system or network, of geographically separated members of a group. The term "teleconference" has been used recently mainly to refer to interactions organized or presided over by or with the aid of programmed computers. In some teleconferences, the members of the group participate concurrently; in others, each member logs in when it is convenient for him to do so, reviews what has happened in his absence, makes his contribution, and logs out, perhaps to return later in the day or later in the week.

During the last five years, a considerable amount of experience has been gained with computer-facilitated teleconferencing, but it is evidently a complex and subtle art, and teleconference programs still have a long way to go before teleconferences approach the naturalness of face-to-face interaction. On the other hand, we note the inefficiency of traveling to meetings and the inefficiency of letting one participant take up the time of $n - 1$ participants when only $m < n - 1$ are interested in what he is saying. As teleconferencing is perfected (especially nonconcurrent teleconferencing), it will become an extremely important technique.

C. Neopaperwork

"Office automation," "computer-based officework," "the high-technology office," and a few other such phrases refer to the aggregation and integration of several applications of computers and networks in office work. ("Automation" is intended in its weak sense, which includes computer "aiding" and "semi-automation.") Office automation includes everything presently called "word processing" (dictation, docun preparation, etc.) plus computer-based filing (informa storage and retrieval), communication (electronic mail, e tronic message services, duologue, teleconferencing), modeling (simulation), and it connects with electronic fu transfer, management information systems, and parts—if all—of computerized commerce.

Office automation is expected to make heavy use of works, both local and geographically distributed. Much of work is organized in an approximately hierarchical man with component desk functions such as transcription, editi filing, retrieval, scheduling, and telephone answering at l echelons and corporate or divisional functions such as planni marketing, operations, and public relations at high echelo Low-echelon functions typically are carried out locally, with a single office or suite of offices, and, when low-echelon fu tions are supported by minicomputers or microcompute local networks will be required in their integration. geographically distributed organizations, of course, geograp cally distributed networks will be required as higher level fun tions are integrated.

1) Telework: Networks will make it possible for people do informational work effectively at locations remote fro their managers, their co-workers, the people who report them, and, indeed, even from customers and clients with who they must interact. Such telework will require facilities fo duologue, teleconferencing, and all the other aspects of offic automation—but little beyond what automation of a nondis persed office will require.

Telework will offer the possibility of saving the hours and the energy spent in commuting. It may burden some families with more togetherness than was contracted for through the marriage vows ("for richer or poorer, but not for lunch"). But its strongest impact on individual lives will surely be felt by persons immobilized by prolonged illness, physical handicap, or children. For many of them, networks will open many doors—including the door to gainful employment.

2) Augmentation of the Intellect: The at-a-distance aspect of computer-based work that is emphasized by the term "telework" will be overshadowed, in the opinion of many, by what Engelbart [18]–[21] has called "augmentation of the intellect." Computers will help people do informational work faster and better by providing fast and accurate tools to supplement such slow and fallible human functions as looking up words in dictionaries, copying references for citation, stepping through checklists, and searching for matching patterns. Augmentation is needed at levels that range from A) helping poor typists who cannot spell to put out neat and accurate reports, to Z) improving the content and style of top-level policy statements. Expectations differ concerning the prospect for significant early contributions from artificial intelligence, but it is clear that relatively unsophisticated augmentation systems can make major contributions. Consider the help provided by descriptor-based and citation-index-based information retrieval systems to a person looking for references pertinent to a particular fact or concept. Or consider the impact that would be made, on writing such as this, by a text editor that automatically displayed the Flesch Count [22] of every paragraph it helped compose.

3) Task Management and Coordination: In addition to helping the individual worker, computers will facilitate teamwork. Each office task will have its planned course of actions, in-

volving particular workers at particular projected times. A computer-based task management system will monitor the task as it moves along the course, checking the actions as they are taken, arranging that planned coordinations and approvals are obtained, and revising the plan (or calling for human help) when the schedule slips. In the early days of office automation, the task management process will be mainly a matter of maintaining orderly work queues for the office workers and displaying for them at each moment 1) what needs to be done and 2) the information needed in doing it. What will need to be done will usually be, of course, to solve a problem or to make a decision—most of the preliminary work will have been performed automatically by computers. With the passage of time, as people come to understand the problem-solving and decision-making processes and the supporting information in programmers' terms, computers will chip away at the problem-solving and decision-making substance of office work, but we expect the now-rising wave of office automation to succeed or fail on the measure of its help to human workers and to human teamwork.

D. Management Applications

Office automation will have its impact upon management, of course, as well as upon the office workers. Management deals almost exclusively with information. (Money is essentially information, of course.) The comptroller's department was computerized early. Electronic funds transfer will be a major application of special-purpose, limited-purpose, or general-purpose networks. On-line financial services may burgeon. Inventory, ordering, production, pricing, and planning will all be interrelated with the aid of networks and computer modeling.

1) Management Information Systems: The widespread feeling of disappointment in the management information systems (MIS's) [23]-[26] of the 1960's and early 1970's had, we believe, a simple basis: the activities that generated the information required to support management decision making had not yet been brought on line to computers, and, therefore, the required information was not available to the management information systems. To some extent, information important to the manager is so global in scope that capturing it all on line is still not possible. (It was not worthwhile to keypunch all the basic operating data just to feed them into the management information system, for only a small fraction of the totality would ever be used. It was impossible to anticipate just what subsets or aggregations of the basic operating data would be required.) As soon as all the informational activities involved in operating an organization are on line, however, the basis for an effective management information system will exist. A few organizations are already approaching that state, but most are just entering—or just beginning to contemplate—office automation.

Local and geographically distributed networks will make it possible, at a cost, for top management to access all the facts and figures involved in the minute-to-minute operations of a business. Top management should resist the temptation to convert that possibility into actuality. The principle that looks best at present is to let the data of a corporation reside where the managers most conversant with them reside ("keep the data near the truth points") and to have conversant managers "sign off" on the release of data upwards in the corporate tree. Certain data should be abstracted and moved upward according to preset schedules; other data may be queried from above—but queried through an authenticating release process. Of course, the release process may in some instances be mediated by programs operating on behalf of the human conversant manager rather than by the human conversant manager himself or herself.

The foregoing discussion pertains, indeed, to most of the data management functions in office automation. In distributed organizations, data will be distributed, and one of the main uses of networks will be to move data from points of residence to points of use.

2) Modeling and Simulation: Computer-based modeling and simulation are applicable to essentially all problem solving and decision making. At present, however, modeling and simulation are computer applications much more than they are network applications—and they are far from ubiquitous even as computer applications.

The trouble at present is that most kinds of modeling and simulation are much more difficult, expensive, and time consuming than intuitive judgment and are cost-effective only under special conditions that can justify and pay for facilities and expertise. But those are prime conditions for resource sharing and, hence, networking. Whereas very large organizations will be able to afford their own concentrations of facilities and expertise, small organizations will not. As management grows tighter and more sophisticated, therefore, there may come to be a place for management consultation and service firms that specialize in modeling and simulation and offer very large or special facilities—and deliver their products through networks. Perhaps a glimpse of such a future has been given by the large array-processing computer, Illiac IV [27], [28], which has been used through the ARPANET in modeling the world climate and the space shuttle. Similarly, MIND [29], a system accessible through a value added packet network, is being used to design communication networks.

E. Commerce

Shifting our attention from activities within an organization, such as a business firm, to interactions among organizations, we can see another kind of application for networks.

1) Electronic Markets: Networks will serve as marketplaces, providing meeting grounds for buyers and sellers. At first, networks will displace telephone and mail, which now serve the marketplace function for most businesses. Later, networks will begin to displace stock exchanges and commodity markets. Ordinary office automation and funds transfer facilities will adequately support negotiations and transactions when the "commodities" bought and sold are purely informational or sufficiently specifiable by words and figures. Wide-band facilities for examining products at the time of purchase ("squeezing the grapefruit") may extend the scope of the electronic marketplace to commodities that must be selected or approved individually by prospective purchasers. We can expect networks to go beyond the role of the mere place or medium for transactions and, with the aid of sophisticated programs, actively to "make a market" in the sense that certain stock brokers make markets in certain stocks.

2) Computerized Commerce: Computerized commerce [30] is based on the idea of electronic markets. It goes beyond providing a marketplace and making a market—and back into the primary motivation of the business firm—by using computers to develop and carry out the strategies and tactics of buying and selling. The concept of computerized commerce is appli-

practice [34], [35]. Human expert consultation will, of course, be available to supplement the computer knowledge bases, but considerations of cost and availability will almost surely favor the computer. Difficult problems of legal responsibility and liability may have to be solved: advice from a knowledge base may be similar to advice from a book, but a knowledge-based program that controls the administration of an anesthetic would appear to introduce a new factor.

Possibly even more far reaching in its implications than access to medical knowledge bases by physicans is access to medical knowledge bases by laymen. Knowledge bases for laymen would have to be quite different in content and packaging from knowledge bases for physicians, and ideally the two applications would complement each other. The layman-oriented application might deal mainly with the complaints not ordinarily taken to a doctor or with the decision process that determines whether or not to seek a physician's help. In either case, if the knowledge-base program had access to the individual's medical record, and if it could make simple observations such as temperature and pulse rate through the network, it could go rather far beyond the limits of the conventional book of medicine for the layman. Society should examine such incursions by the computer into medicine or even paramedicine very carefully before making up its collective mind about them. They obviously mix benefits with dangers. Unfortunately, they tend to be approached with prejudice.

G. Government Applications

Actual and potential government applications of networks include military command and control, communications, logistics, acquisition and interpretation of intelligence data, dissemination of intelligence, law enforcement, delivery of government services such as Social Security benefits to citizens, and converting the paperwork of the bureaucracy into bits. Paperwork in the government is rather like paperwork in the private sector, but carried a step or two further into detail. The other government applications, on the other hand, seem rather special. Military command and control, communications, intelligence, and to a considerable extent logistics systems must be able to operate fast, move fast or hide, and function in the presence of physical (as well as other) countermeasures. Law enforcement information systems are in some ways like highly amplified credit reference systems: derrogatory information seems especially crucial, for to be forewarned is to be forearmed, and action must often be taken on the basis of whatever data can be assembled in a few seconds. Serving all the citizens and collecting taxes from most of them requires that certain personal data be held about almost everyone—enough in sum to make a several-trillion-character data base that at least conceivably could be subverted to political or economic exploitation. There are strong lessons about government applications of networks in the recent rejection by the Office of Management and the Budget (OMB), at the well-timed suggestion of several members of Congress, of the proposed new Tax Administration System of the Internal Revenue Service.

1) Military Command and Control and Military Communications: Military command and control and military communications are prime network applications. Both interactive computing and networking had their origins in the SAGE system (Semi-Automatic Ground Environment for air defense), and many of the military systems used to command forces and [illegible] works. For reasons we do not fully understand—since fast response to a changing situation is the essence of command and control—the World-Wide Military Command and Control System (WWMCCS) is actually not very interactive, and its computers, which use the GECOS operating system [36] designed for batch processing, are not interconnected by an electronic network. But surely WWMCCS will in due course be upgraded. Autodin II is under development and will supplement or replace Autodin I [37], the Department of Defense's present store-and-forward digital telecommunications network, with a modern packet-switching network based on modified and secured ARPANET technology. Networking is being pursued actively, also, in the intelligence community. One of the most significant possibilities for the military that is opened up by advances in information technology is the achievement of a much tighter coupling between intelligence and command and control. One can envision a reduction in the time required for the distribution of intelligence information from days or hours to minutes or seconds. Such an advance would, of course, put pressure on intelligence gathering and processing to operate on faster time scales.

2) Military Logistics: There is less progress, but also less pressure, in the logistics area, where more than 20 large batch inventory systems can be counted, diverse in respect of hardware, programming language, and data management system. Over the coming years, however, even the logistics situation will probably be brought under control and onto a network. The overall objective is to make the entire operation of a military effort responsive to coherent hierarchical command in the light of valid and current intelligence—with security against enemy actions and countermeasures.

3) The Network of the National Crime Information Center (NCIC): The NCIC is operated by the FBI and connects with state and local police units in most of the states. The NCIC contains, among other things, data on stolen cars and stolen license plates and the police histories of convicted criminals. The case of the NCIC network is an interesting study because, in it, the informational needs of the police and the information-providing capabilities of computers and telecommunications run head-on into Congressional concern for the right of informational privacy. When a police officer stops a speeding car and approaches it to make an arrest, he would like to know something about the car and driver. Is the car stolen? Does the owner have a history of resisting arrest? Forewarned is forearmed. About two years ago, however, an innocent man was killed by an arresting officer forearmed by forewarning with incorrect information. In the most recent chapter, the Senate Committee on Government Operations caused to be rejected the NCIC's request for permission to acquire a message-switched network to speed up communication with state and local police.

4) Social Security: In the U.S., the Social Security Administration (SSA) distributes more than $100 billion a year to more than 20 million people and interacts with millions of clients each year through about 1300 offices manned full-time and 3000 manned part-time. The set of computer processible data bases that support SSA operations contains more than a trillion characters, and it is estimated that in those operations each year several trillions of characters flow from one location to another, about a tenth of a trillion by a network and processing system called the "SSADARS system," and the rest mainly [illegible]

In 1976, the SSA began planning the modernization of the process through which it discharges its massive responsibilities [38]. The new process will make even heavier use of computers than does the present one and there will be much consultation and updating of central or regional data bases from local offices. (The present process requires several computer areas, each with multiple mainframes, more than a hundred disk drives, and more than a hundred tape drives—and, in all, approximately 400 000 magnetic tapes.) The new communication subsystem will, therefore, be a network of very major proportions, probably a dedicated SSA network operated by the General Services Administration or (improbably in the near term) a part of an even larger and more comprehensive network. Most technologically developed countries will sooner or later have social security networks.

H. Protection

If military, intelligence, and police networks are reckoned as networks for protection, then protection is a very large category of network applications. There is another member of the class that deserves mention.

1) Home and Neighborhood Security: Several of the projected applications of computers in the home relate to security: sentry against intrusion, fire, and gas and water leakage, monitoring the well-being of the elderly and infirm, and "electronic babysitting." In most of these applications, computers will be better at detecting trouble than in correcting it, and there will be a strong requirement for communication with remote persons or agencies. At present, some burglar alarms are connected by dedicated lines to central security offices or police stations and some "dial up" in the event of trouble. If a packet network were available, it would probably be less expensive and it would provide a wider range of options, including absent members of the family, friends, and neighbors as well as security companies and public agencies.

It seems possible that a neighborhood communication medium (with a broader fan-out or faster sequencing of calls than the telephone) might be just what is required for the elderly to help one another achieve a higher level of security and peace of mind. CB radio or house-to-house (or apartment-to-apartment) wiring or a multipurpose packet network could provide the medium.

Probably just conversation of the kind that prevails on CB radio interconnections would go a long way, but it could be reinforced by slightly higher technology. Home computers could be programmed to interpret a variety of indicators of trouble—the sound of a fall, too long a flow of water, the refrigerator door open, prolonged quiescence—and to ask for an "all's well" report whenever there was cause for concern. Failing to be satisfied that all was indeed well, the computer could call for help. It would have a list of participating neighbors and a schedule of probable availability for each, and—by communicating with their home computers—it could quickly find someone to look in and check, or provide assistance. It has been suggested that a neighborhood net could monitor its clients while they were walking on the sidewalks as well as while they were at home, and a small device has been demonstrated that sends out a radio signal when its wearer falls down—or for some other reason becomes horizontal [31]. If neighborhood networks existed, there would probably be no end of inventions to exploit them in the interest of security: heartbeat monitors, breathing monitors, footstep sensors, and so on. And if the present trend of population statistics continues, security applications might constitute a significant sector of the network application pie.

I. Education and Awareness

Beginning with the last section and continuing now into this one, the focus of interest has moved from the organization—or the individual as a member of an organization—to the individual as an individual in the primary family group in the home. Probably the most important network prospects for the individual in this century lie in education and training.

1) Computer-Based Education and Training: We assume that advances in computer representation of knowledge and in computer mediation of interactions between people and knowledge bases will advance computer-based education and training far beyond the "expensive page turners" and drill and practice routines that are associated in many minds with the term "computer assisted instruction." We assume that knowledge bases accessible through networks will eventually accumulate more knowledge, in each of many fields of learning, than typical teachers are able to master and retain, and that the knowledge in the knowledge bases will be well organized (by experts in each field) and effectively accessible to students at all levels of mastery and aptitude. However, computer-based techniques for the representation, organization, and exploration of knowledge are at present still topics of research—and even if they were fully developed today, it would still take a decade or two to translate the content of the many fields of learning into computer-processible knowledge bases. During the coming years, therefore, application of networks in the area of computer-based education and training will be preliminary and propaedeutic. Perhaps toward the end of the century it will approach its ultimate volume and significance—and be among the top three or four uses of networks.

2) News: At present, most people gain their awareness of what is going on in the world mainly through mass media that report on events rather than processes, that select a few news items instead of covering the news, and that give everyone, regardless of his or her interest pattern, the same few selections. Networking has the potential of changing the news into a multidimensional dynamic model of the world that each individual can explore in his own way, selecting for himself the topics, the time scales, the levels of depth and detail, and the modes of interrogation and presentation. Interest profiles and other techniques of selective dissemination may play important roles, but networking in principle removes the necessity of disseminating (with its implication that the initiative lies mainly with the transmitter) and opens the door to self-directed exploration and investigation by the receiver of the news. To provide the multidimensional dynamic model for exploration and investigation would, of course, be a demanding responsibility for the gatherers and organizers of news, but they gather and organize much more even now than they print or (especially) broadcast. There will probably be a long slow evolution from the newspaper/newsmagazine format and the nightly news format through increasing levels of user initiative toward truly user-dominated interaction with a whole-world knowledge base.

III. Requirements Imposed Upon Networks by Applications

Now that we have sketched out several applications, we should examine briefly the network characteristics they require. The applications do not all require the same network characteristics, of course. One application may require one

pattern of characteristics, while another application may require another pattern. Some of the frequently required characteristics are the following.

1) Bidirectional Transmission: Most applications require two-way communication—if not the capability of sending and receiving simultaneously (full duplex), then at least the capability of alternating between sending and receiving (half duplex).

2) Freedom from Error: One wrong bit may completely change the meaning, especially if numerical data are represented nonredundantly. In such cases, even though the basic communication channel itself is not error free, the end-to-end communications must be made error free through the use of adequate error handling mechanisms.

3) Efficiency Despite Burstiness: A source that transmits short bursts of information and is quiescent between bursts typically does not wish to pay for channel time while it is quiescent. Both human beings and computers are bursty sources.

4) Low Cost per Bit: The cost of network service depends, of course, upon many nontechnical factors as well as upon the technical efficiency of the network in converting its resources into services. But technical efficiency is a very strong and basic factor. This characteristic refers to the cost of transmitting one bit from source to destination. The relation of cost to distance is considered separately.

5) High Connectivity: A source may need to transmit to any one or more of many destinations. A destination (i.e., user) may need to examine many sources.

6) High Information Rate: Wide-band channels are capable of transmitting many bits per second. The criteria for "high," "wide," and "many" vary widely with type of signal and level of expectation. The 50 000-bits/s information rate of most of the ARPANET channels seems like a high information rate for ordinary interactive computing, but it is too low for convenient transmission of large files (e.g., high-resolution photographs) and far too low for moving pictures, even low-resolution television pictures.

7) Security: Security is a complex of characteristics, some of which provide the technical basis for the protection of privacy. Others have to do with preventing disruption of service and protecting against fraud and theft.

8) Privacy: In the U.S., this complex of informational rights, including but by no means limited to protection against eavesdropping, has been formulated by the President's Commission on Privacy and, to a considerable extent, expressed in legislation in the Privacy Act of 1977. Other nations also have privacy laws, of course, some of them in some respects more stringent than ours.

9) Authentication: A good authentication scheme provides the electronic equivalent of a signature. Ideally, authentication identifies the author of a document and makes it impossible for him to escape responsibility for the authorship. Ideally, also, authentication makes it impossible for anyone to change even one character or bit of the document without destroying the "signature."

10) High Reliability: Low probability that network service, as seen by the application, will be impaired by macroscopic malfunctions. For present purposes, we distinguish between macroscopic malfunctions and microscopic errors in bit transmission.

11) Full-Duplex Transmission: Some applications require, and most are favored by, the capability of sending to another station and receiving from it at the same time.

12) Priority Service: Guaranteed or preferential service, especially when the network is congested, is widely regarded as essential for certain very important functions or for certain very important persons.

13) Speech Capability: Present speech circuits transmit alphanumeric information inefficiently, and most present data networks were not designed to transmit speech. It will be advantageous, however, to integrate speech with data.

14) Pictures: It will be advantageous to integrate pictures, also, into the repertoire. Graphs, charts, diagrams, and simple sketches fit readily into the pattern of data transmission, but high-resolution pictures and, especially, moving pictures require high information transmission rates. This characteristic is essentially a second "information rate"—but scaled in such a way as to be more demanding of very-wide-band capabilities.

15) Insensitivity to Distance: Synchronous satellites and packet switching both tend to make the difficulty and cost of transmission less dependent on distance than they are in traditional communication systems. Rarely is it an absolute requirement that difficulty and cost be independent of distance, but often it is desirable.

16) Short Transit Time Delay: If the sum of the signal-transit time and the signal-waiting-in-buffer time is too great, an application may be slowed down too much or disrupted. The 0.2-s delay introduced by transmission via a synchronous satellite somewhat disturbs two-way speech communication. The delay introduced by transmission from one processor to another may slow down the operation of a multiprocessor that is a network of minicomputers or microcomputers.

17) Uniform Time Delay: In some applications, successive segments of the signal must reach the destination in sequence (or be put back into sequence if they arrive in scrambled order). Note that reordering may cause all the segments to be delayed as much as the most-delayed segment.

18) Broadcast Capability: Some applications require, and some are favored by, the capability of transmitting to many or all destinations concurrently.

19) Mobility: Some or all of the stations may need to move from place to place and may need to communicate in transit.

To obtain rough measures of the requirements imposed upon networks by the several applications, we filled in the body of (an early version of) Table I.[1] Into each cell we entered a number to indicate our intuitive rating of the importance of the characteristic for the application. The rating scale we used runs from 0 (lowest) to 5 (highest). For example, we considered connectivity to rate at 4 in importance for mail and message systems because mail and messages typically fan out widely from senders to receivers and fan in to receivers from a wide distribution of senders. We did not assign a 5 because mail and message systems would still be valuable (cf., the plans of Satellite Business Systems) if connectivity were limited to within organizations. In the case of the column 6, information rate, we used a somewhat special scheme. The numbers from 1 through 5 encode five class intervals of information rate in

[1] In order to obtain a broader basis on which to think about, and possibly model, the relations between networks and applications, we suggest that you (the reader) photocopy Table I and fill in some rows with your estimates of the importance of the characteristics to the applications. We have also provided room for you to define additional applications or characteristics. If you are willing to share your estimates with us, despite the fact we are not bold enough to share our raw-data estimates with you, please post them to J. C. R. Licklider and A. Vezza, MIT-LCS Rm. 219, 545 Tech Sq., Cambridge, MA 02139.

TABLE I

IMPORTANCE COEFFICIENTS OF NETWORK CHARACTERISTICS FOR VARIOUS CLASSES OF NETWORK APPLICATIONS (EACH CELL ENTRY SHOULD BE SUBJECTIVE RATING OF IMPORTANCE ON A SCALE FROM 0 (LOW) TO 5 (HIGH))

NETWORK CHARACTERISTICS / NETWORK APPLICATIONS	1 Bi-Directionality	2 Freedom from Error	3 Efficiency Despite Burstiness	4 Low Cost Per Bit	5 Connectivity	6 Information Rate	7 Security	8 Privacy	9 Authentication	10 Reliability	11 Full-Duplex	12 Priority	13 Speech Capability	14 Picture Capability	15 Insensitivity to Distance	16 Shortness of Delay	17 Uniformity of Delay	18 Broadcast Capability	19 Mobility	SPACE FOR ADDITIONAL CHARACTERISTICS						SUBJECTIVE RATING OF IMPORTANCE OF APPLICATION
BASIC																										
[illegible]																										
Storage																										
[illegible]																										
Information																										
COMMUNICATION																										
[illegible]																										
Duologue																										
[illegible]																										
Speech																										
[illegible]																										
Still Pictures																										
[illegible]																										
NEOPAPERWORK																										
[illegible]																										
Augmentation																										
[illegible]																										
MANAGEMENT																										
[illegible]																										
Modeling																										
COMMERCE																										
Electronic Markets																										
Computerized Commerce																										
[illegible]																										
PROFESSIONAL																										
Monitoring Patients																										
Monitoring Non-Patients																										
Medical Records																										
Medical Knowledge Bases																										
GOVERNMENT																										
[illegible]																										
Logistics																										
[illegible]																										
Social Security																										
PROTECTION																										
Home Security																										
EDUCATION AND AWARENESS																										
Computer Based Education																										
News																										
SPACE FOR ADDITIONAL APPLICATIONS																										

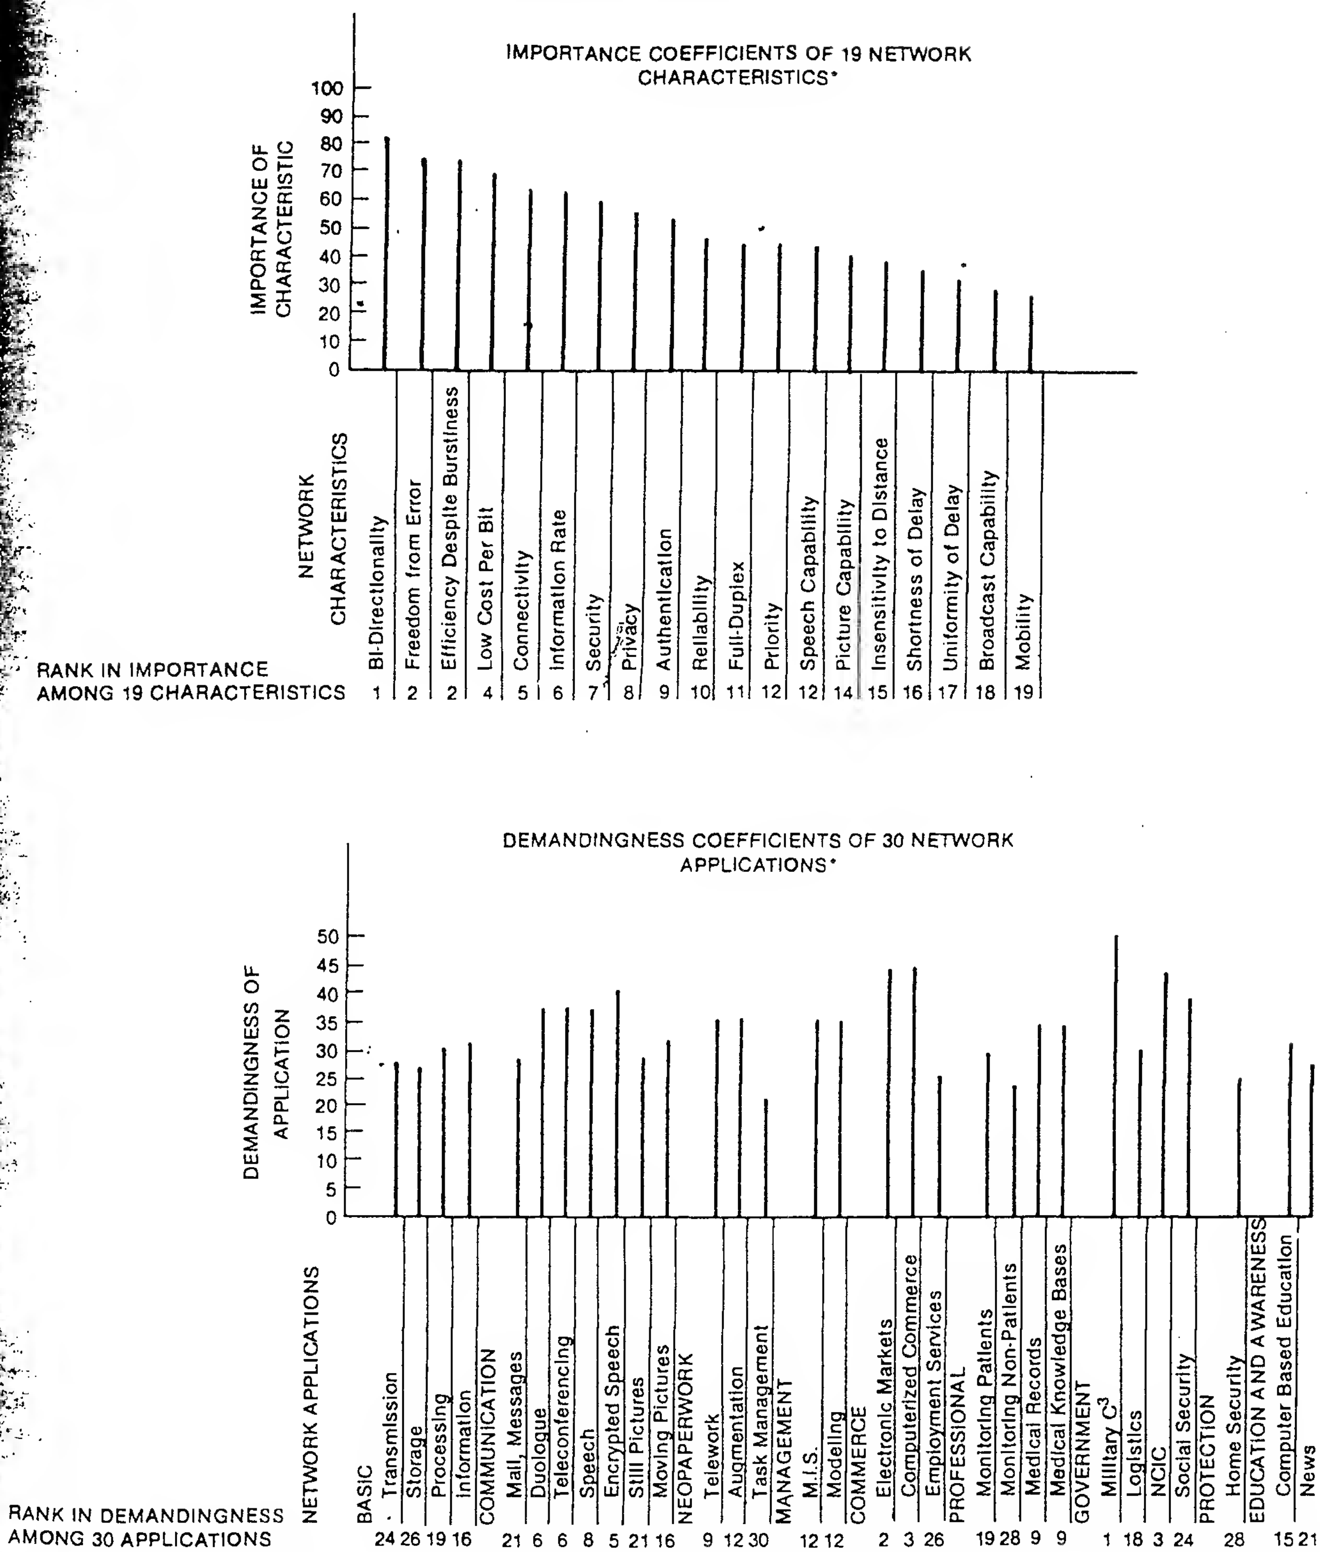

Fig. 1. Importance Coefficients of 19 network characteristics and demandingness coefficients of 30 network applications. (*Divide each number by 1000 to scale the sum to unity.)

bits/s: 1) 75–300, 2) 300–1000, 3) 1000–10 000, 4) 10 000–1 000 000, and 5) above 1 000 000. For the rough purposes of our analysis, nevertheless, we shall interpret the entries in column 6, as all the other entries, to be estimates of the importance of the (columnar) characteristic for the (row) application.

Fig. 1 shows the relative importance of the 19 network characteristics. As one examines the average ratings of the characteristics, it comes as no surprise that bidirectionality is very important. It is the "co" in "communication."

It may be slightly surprising, however, that freedom from error is so important. It is freedom from error as seen by the

application, of course. There are bound to be errors in the raw network channels, but they may be detected and eliminated by error-correcting circuits or by retransmission. "Error-free" may in practice mean one bit error in 10^{12} or 10^{14} bits, on the average. The importance of achieving an extremely low error rate stems in part from the fact that many of the applications involve information, such as financial data, in which changing a single character could make a great difference. Freedom from error is required, also, by most cryptographic schemes. Where freedom from error is not required by an application, one can usually find error detecting and correcting mechanisms within the application itself. Such mechanisms are quite evident, for example, in human conversation. But it greatly simplifies most network applications if the network can be counted on to do the error handling.

Ability to handle bursty transmissions efficiently ranks third. The advantage provided by this characteristic translates directly into a cost advantage.

Low cost ranks fourth in importance. It did not rank higher because we recognized that certain of the applications, such as military command, control, and communication, are relatively insensitive to cost. Also, other network applications such as mail and messages are already quite cost competitive with their conventional counterparts and do not demand very-low-cost facilities.

Connectivity ranks fifth in importance. The reason connectivity does not rank higher is that we assigned only a medium score for connectivity to applications that required only connectivity within an organization or within a region, and many applications could function—though perhaps at some disadvantage—with such limited connectivity.

Information rate ranks sixth. We interpret that to mean that very wide-band transmission is not vital to most of the applications and that most of them could be satisfied with an information rate in the range 1000–10 000 bits/s. However, that is the information rate seen by the application. To handle heavy traffic, and to handle a few of the applications, a network should have channels of considerably greater bandwidth than that.

Security, the complex that includes assurance of service when required and protection against fraud and theft, ranks seventh.

Privacy, the complex that includes protection against disclosure of personal information and unauthorized use of it, ranks eighth.

Authentication, a characteristic closely associated with security, ranks ninth.

At the other end of the ranking, mobility (19th) is not required by most of the classes of applications we considered—but, of course, is essential for some applications.

Broadcast capability (18th) was scored low because it is not needed at all in many applications and is needed only occasionally in others such as mail and message systems—and, when needed, usually can be simulated adequately by repeated point-to-point transmissions.

Uniformity of time delay (17th) is important mainly for speech transmission. If speech had been given a weighting proportional to its probable eventual importance in networking, uniformity of time delay would have ranked higher.

The capability of giving preferential treatment to high priority traffic (12th) ranks as high as it does because we viewed priority in the context of present-day systems that may introduce considerable delays into the delivery of some or all of their messages. In the context of future systems in which a whole transmission will take less than a second, priority may be much less important. However, the need for priority is unlikely to vanish. Priority classes are useful in queuing messages for processing by people and in indicating the prioritizer's sense of urgency or importance to the recipient. Moreover, even very wide-band systems tend to be designed just barely to handle expected peak loads, and even such systems can be overloaded—in which case, prioritization might be helpful. On the other hand, it is conceivable that the processing of priorities might slow a system down more than eliminating low priority traffic could speed it up.

We do not want to attribute too much value to our no-doubt-ideosyncratic subjective estimates of the importance of network characteristics to applications, but we would like to carry the analysis another step to illustrate what we think might be a valuable method. It attempts to deal with the relative merits of various networks or network architectures.

The method begins with a table of applications versus characteristics similar to Table I except that each application has an importance weight and all its cell values are multiplied by that weight before the columns are totaled. The method assumes, also, a table of networks or network architectures versus network characteristics such as Table II. The entries in Table II represent our very subjective impressions of the degrees to which the characteristics at the top characterize the networks at the left-hand side. The values are certainly not definitive. In the case of the hypothetical augmentation of the ARPANET, they assume major increases in number of subscribers and in information rate, and they assume that advanced provisions are made for security, privacy, authentication, and priority service. They assume, also, that satellite relays are incorporated into the network along with wide-band surface channels and that there is a packet-radio subsystem to serve mobile applications.

In the case of the projected SBS service, indeed, they are based only on the most informal information, and they make rather optimistic estimates about the characteristics of the hypothetical networks that would be developed on the basis of the SBS facilities. The reader is invited to substitute his or her own estimates.

To determine the suitability of a network or network architecture to a set of applications, one simply multiplies each cell value in its row in (the table like) Table II by the importance of the corresponding characteristic at the bottom of (the table like) Table I—and then finds the sum (across the row) of the products.

To illustrate the use of the method, we worked with the four networks of Table II and with four sets of applications. Application set 1 was the set shown in Table I. Sets 2, 3, and 4 were subsets consisting of—set 2: speech and encrypted speech, set 3: still and moving pictures, and set 4: duologue and augmentation. (We gave equal or uniform weighting to the applications in each set—to all 30 in the first set and to both of the two in each of sets 2, 3, and 4.)

We obtained the 16 appropriateness indexes shown in Table III. The dial telephone network performs best in the speech applications, of course, but it does not appear to do badly in the others. (Giving more weight to cost tends to reduce its scores.) The experimental ARPANET appears to perform well on speech, which surprised us despite the fact that experiments on the transmission of compressed speech over the ARPANET have been very successful, but best on duologue and augmentation, which we expected. In the scoring, the ARPANET suffers because, being experimental, it was not developed in respect to some of the important network characteristics. Assuming such

TABLE II
ESTIMATED DEGREES TO WHICH FOUR SELECTED NETWORKS POSSESS THE 19 NETWORK CHARACTERISTICS (THE RATINGS ARE THE AUTHORS' INTUITIVE ESTIMATES ON A SCALE FROM 0 (LOW) TO 5 (HIGH)

NETWORK CHARACTERISTICS / NETWORKS	1 Bi-Directionality	2 Freedom from Error	3 Efficiency Despite Burstiness	4 Low Cost Per Bit	5 Connectivity	6 Information Rate	7 Security	8 Privacy	9 Authentication	10 Reliability	11 Full-Duplex	12 Priority	13 Speech Capability	14 Picture Capability	15 Insensitivity to Distance	16 Shortness of Delay	17 Uniformity of Delay	18 Broadcast Capability	19 Mobility
Dial Telephone	5	2	0	2	5	3	2	2	1	3	2	1	4	2	0	4	4	2	2
ARPANET	5	4	4	3	3	3	1	2	1	4	4	0	2	2	4	3	3	1	1
Hypothetical Augmented ARPANET*	5	4	4	4	5	5	4	4	4	4	4	3	4	4	4	4	4	3	4
Hypothetical Corporate Network Using SBS**	5	4	4	4	2	5	4	4	4	4	4	3	4	5	4	2	4	3	0

*A hypothetical network based on ARPANET technology but with very wide-band ground and satellite channels, very many subscribers, advanced provisions for security, an authentication scheme, arrangements for priority, and a mobile/portable radio adjunct based on the ARPA Radio Net.

**A hypothetical network of the kind that might be based on the projected facilities of the Satellite Business Systems Corporation and used by a large corporation with geographically distributed branches. It is assumed that this network is used only within the corporation and therefore has restricted connectivity.

TABLE III
APPROPRIATENESS SCORES FOR FOUR SELECTED NETWORKS, EACH RATED ON FOUR DIFFERENT SETS OF APPLICATIONS (THE APPROPRIATENESS SCORES ARE BASED ON A SCALE FROM 0 (LOW) TO 5 (HIGH). THE WAY THEY WERE DETERMINED IS DESCRIBED IN THE TEXT.)

APPLICATION SETS / NETWORKS	1 30 Applications of Table I	2 Speech and Encrypted Speech	3 Still and Moving Pictures	4 Duologue and Augmentation
Dial Telephone	2.5	2.8	2.4	2.6
Experimental ARPANET	2.8	2.8	2.5	3.0
Hypothetical Augmented ARPANET	4.1	3.9	4.0	3.8
Hypothetical Corporate Net Based on SBS Service	3.8	3.5	3.9	3.8

development (Hypothetical Augmented ARPANET) yields the appropriateness scores in the third row of Table III, which are all toward the upper end of the five-point scale. The hypothetical corporate network based on the projected SBS service appears to perform almost as well as the hypothetically augmented ARPANET, suffering in the comparison only because we assumed for it limited connectivity (intercorporate communication only), no surface channels (and therefore always the 0.2-s satellite delay), and no mobility. Those lacks showed up only in the average over the 30 applications and in the speech applications.

Obviously, the result obtained with the method is no better than the ratings it processes, and we do not make any claims for our ratings. We believe, nonetheless, that the scheme puts into an orderly array some of the basic factors that determine the relative appropriatenesses of various networks for various sets of applications, and that it leads the users of the method—or at least it led us, as we used it—to consider the factors carefully and to think about how they act and interact. Upon examining the interactions, it soon becomes clear that the linear weighting scheme smooths over many nonlinear logical interactions, and that a more advanced model would have to be more like a computer program than three tables and a pocket calculator. Nevertheless, the first step has to be to survey the variables that are active in networking, and the simple scheme provides a start on that.

IV. ISSUES

A. Brittleness

Brittleness is approximately the inverse of the lauded complex of system attributes: flexibility, robustness, and gracefulness of degradation. Brittleness often arises from a quest for efficiency or economy. If you space the pony express posts as far apart as a fresh pony can run, then the mail does not get through when an emergency forces you to use a tired pony. A socioeconomic unit with a minimum-capacitance supply system avoids the waste inherent in having products stagnate in the pipeline but crumbles in a siege. Networks will almost surely be more efficient than the systems they supplant. Should not that expectation prompt us to ask whether they will also be more brittle?

From an engineering point of view, the preferred approach is to avoid brittleness through judicious choices in the architecture and design of networks. The dynamic routing feature

of the ARPANET, for example, permits the network to continue to operate, as far as two functionally connected host computers are concerned, as long as there is some path left intact between them. Indeed, just as it is to reliability, redundancy is the main engineering antidote to brittleness in networks—redundancy of interconnection, redundancy of power supply, and redundancy of information storage. Moreover, sophisticated uses of redundancy, as in restructurable logic and error-detecting codes, provide much greater returns in robustness than do brute force applications that cost the same.

B. Electronic Imperialism or Technology Transfer to the World

Aviation may have had more impact in technologically not-yet-developed but developing countries with poor roads and few rails than in technologically developed countries that already had working transportation systems before airmail routes and airlines came upon the scene. Brazilia, for example, would not have been feasible without the DC-3. Analogously, networks may have their most dramatic effects where there are few critical masses of knowledge and few self-reinforcing centers of intellectual activity. Networks may link the geographically separated subcritical foci of cognition in the developing world with the concentrated supercritical centers of the developed world, bringing the former deeply into the interaction patterns of the latter and making it much easier for the former to grow and advance. If networks do turn out to have such an effect, it may represent a new dimension of imperialism, or it may open up broad new avenues of technology transfer, or, as seems most likely, it may look one way to some and the other way to others.

Although technology transfer to the developing world is often viewed as a matter of delivering journals, reports, and books and of transmitting data to third-world countries, truly effective transfer is a transfer not of data or of information but of knowledge and it flows through human interaction. Perhaps the most effective pattern involves graduate study by promising young people from developing countries in leading graduate schools in developed countries—followed by work experience in the developed countries and then return to form foci of advanced technology in the developing countries. But the difficulties with even that pattern are well known: reluctance to return resulting in the "brain drain," or isolation after return resulting in unhappiness and ineffectuality. What is needed is a way to return without breaking the ties of interaction with teachers, fellow students, managers, and co-workers in the centers of technology that, thanks to their highly developed intellectual and motivational supports, made possible the transfer of knowledge in the first place. Networks can fill that need. The ARPANET has provided several instances, albeit just in the U.S., in which students have remained functionally and motivationally in the research groups in which they worked for their degrees until they could build up self-sustaining foci of their own in their postdoctoral locations. It seems very likely that the postdoctoral locations could as well be in foreign countries, even countries with little technology, if only they had network connections and sufficient funding to keep local terminals in reliable operation and to pay for computing time and storage somewhere on the net. Indeed, it seems likely that technology transfer to developing countries could become real and effective through no more than informal extension of patterns of interaction that have become well established in the ARPANET community. But there is no reason to keep the patterns entirely informal. Formal associations between universities, not-for-profit organizations, and business firms in the have-technology and the have-not-technology countries would surely increase the productiveness of the technology-transfer enterprise [39].

We do not want to try to take sides with respect to, and we cannot hope to offer a solution to the problem suggested by the juxtaposition of "electronic imperialism" and "technology transfer." We do suggest, however, that the prospect of networking between the developed and the developing worlds deserves very serious study. From the point of view of the imperialist, it may well be that packet networks will be to the not-far-distant future what clipper ships (or were they packet ships?) and clipper aircraft were to the not-far-distant past. From the point of view of the countries that need and want technology transfer, packet links to technologically developed countries may be by far the best way to get it if they can figure out how to keep the electronic colonialism from coming with it.

C. Unity, Federation, or Fragmentation

If we could look in on the future at, say, the year 2000, would we see a unity, a federation, or a fragmentation? That is: would we see a single multipurpose network encompassing all applications and serving everyone? Or a more or less coherent system of intercommunicating networks? Or an incoherent assortment of isolated noncommunicating networks, most of them dedicated to single functions or serving single organizations? The first alternative—the strongly unified network—seems improbable: of almost zero probability if the scope is taken as world-wide and still of very low probability even if the scope is taken to be the U.S., which seems to engender pluralistic solutions to most problems. The third alternative—many separate noninterconnecting networks—is what would be reached by proceeding with a plan and, therefore, may be judged rather probable, but it would be very disappointing to all those who hope that the whole will be much greater than the sum of the parts, i.e., that many of the projected applications will facilitate and contribute to one another to such an extent that the overall value will grow combinatorially. The middle alternative—the more or less coherent network of networks—appears to have a fair probability and also to be desirable, but it brings with it the problem of how to achieve enough coherence to support fast and facile intercommunications among the subnetworks when required, and that may be a difficult problem. Let us consider first why coherence is desirable and then turn to the difficulty of achieving it.

"Coherence" characterizes a system in which all the parts articulate well and function in synergy and in which the subsystems are compatible and cooperative. In an information network, coherence is desirable partly for the same reason it is desirable in a telephone system: the value to a typical user increases as the number of other accessible users increases. This is true no matter whether the users are people or computers. The importance of coherence is amplified, however, by the fact that some of the applications of computer networks will involve several functions operating on common information. Planning a trip, for example, will require interaction among: your calendar program and the calendar programs of people you will visit, the reservations programs of airlines, car rental services, and hotels, the funds transfer systems of your bank and several other banks, your company's

office, and perhaps data bases pertinent to business to be transacted on the trip. The computers could not be of much help in the planning if each function or service had its separate network or if their networks were physically interconnectable but incompatible at various levels of protocol. Indeed, coherent interconnection of diverse functions will be essential if networks are to live up to expectations in electronic message services, computerized commerce, delivery of social services that involve both federal and state governments, and many military and intelligence applications.

Coherence, however, is a condition that has to be planned and striven for. It does not arise in a short time through evolution—at least not through evolution that conforms to the spontaneous-variation-plus-natural-selection model. Will the forces that are operating in the present network situation foster a sufficient degree of coherence for networks to fulfill their promise? On the positive side is the fact that a standard packet-switching interface protocol X.25 [40] was formulated and agreed upon in an unusually short time and that the interconnection of such dissimilar networks as Telenet [3] and the Canadian data network has already occurred. Also on the positive side is the possibility that one or two commercial networks, such as the one being developed by AT&T and the one being developed by Satellite Business Systems, will dominate the network market and thereby create the kind of coherence that IBM has created in a large part of the computer software field.

However, the other factors seem to work against coherent interconnection. First, the network situation is evolving without any national policy. (Within the U.S. government, it is evolving without any federal policy.) Several countries are building networks independently. Many companies are building networks independently. Second, although the lines may be leased from the same telephone company that leases lines to everyone else, there is some security in having one's own dedicated network. The need for security in such areas as EFT may be stronger than the need for interconnection. Third, wherever personal information is concerned, and especially in the federal government, privacy has become a major issue, and a simplistic interpretation of the privacy problem sets interconnection into opposition with privacy. Fourth, research and development in the area of network security and privacy assurance have not been and are not being supported at a high enough level to create—soon enough—a technology that will let one say: "You can have both interconnection and security, both interconnection and privacy; you can have your cake and eat it, too." Actually, the technology of communications security is rather well developed [41]–[44]—except for uncertainties arising from the "56-bit controversy" [45]—as a result of many years of work in the military intelligence area, but the technology of computer security is less well developed [46], and nontechnological aspects—plant, personnel, and operational aspects—of network security are not in good condition at all.

The conclusion with respect to unification, federation, or fragmentation must be: we should strive for the kind of federation of networks that will provide coherent interconnection where needed and justified and, at the same time, provide informational privacy and security. That will require planning at national and international levels. It will require intensified research and development in network security and in internetting. And it will require an elevation of the ongoing discourse about privacy—to a level on which legislative and administrative policies can be defined clearly and networks can be designed responsively.

D. Privacy

It is now very widely understood that the collection of large amounts of personal information in computer processible data banks tends to jeopardize personal privacy. The main reason, of course, is that aggregation and computerization open up the possibility of invading privacy efficiently and on a massive scale. At the same time, they open up the possibility of protecting privacy tirelessly and algorithmically and of using the personal data effectively in the effort to accomplish the legitimate purposes for which the data were collected. Indeed, the stage is set for a battle between the forces of good and evil.

The stage is, however, not set in reasonable balance. The things required for the protection of privacy in information networks are policy and technology: legislative and administrative policy to define what is to be achieved and a technological basis for achieving it. In the U.S., the legislation is the Privacy Act of 1974 and the administrative policy is an OMB Circular [47]. Both are cast in terms of absolutes. The technological basis, as mentioned in the preceding section, is a combination of communications security and computer security. Because computer security is a relatively new and neglected subject, it is difficult to provide convincing assurance to an intelligent skeptic that any proposed interconnection of personal data, transmission channels, information processors, and interrogation-and-display facilities will not jeopardize privacy. Repeatedly, indeed, the advocates of proposed federal data networks have failed to present convincing analyses of the threats to privacy and of the trade-offs between privacy and mission effectiveness—and, repeatedly, their requests for permission to procure such networks (e.g., FEDNET, NCIC upgrade, IRS Tax Administrative System) have been sidetracked or denied with good reason.

At least in government circles, therefore, the issue of privacy is a very real and central network issue. Before it can be solved, three things have to be done. 1) The technology of information security has to be improved to the point at which reasonable analyses can be made and assurances can be given. 2) Network advocates have to develop plans and justifications that take privacy into account and provide strong assurances that it will be protected. And 3) the members of the oversight committees and their staffs have to face the fact that to protect privacy by precluding interconnection is not a very satisfactory solution for the long term. They should support and foster the accomplishment of steps 1) and 2) with the aim of receiving plans and proposals that they would not have to kill.

E. Other Issues

Space limitation precludes substantial discussion of other issues, but there are several that should be discussed. We shall discuss only a few of them, and those only very briefly.

1) Transborder Data Flow: Several European countries are beginning to restrict the flow of personal (or personnel) data across their borders on the ground that they must protect the informational privacy of their citizens against threats implicit in data processing in countries with less stringent privacy laws. Some of the countries have laws or regulations that preclude the transmission of encrypted information through their public communication facilities. Many believe that such restrictions may be used to discriminate against foreign (e.g.,

American) data processing firms and against multinational corporations.

2) Technology Export and Import: International networks can be expected to facilitate greatly the transfer of scientific and technical information and know-how among the technologically developed nations. From a nationalistic point of view, one can see both advantages and disadvantages in such transfer to ideological, military, and economic competitors. The advantages are mainly humanistic and short-term economic. The disadvantages are mainly security-related and long-term economic. The interplay of advantages and disadvantages is giving rise to issues that will probably intensify.

3) Competition Versus Monopoly and Free Enterprise Versus Regulation: These are the issues of the "Bell Bill" (Consumer Communications Reform Act), Computer Inquiry II, and several recent decisions of the Federal Communications Commission that have favored competition in the telecommunications industry. How these issues are settled will to a large extent determine who operates the networks of the future in the U.S. and how such applications as electronic message service and "the office of the future" are implemented.

4) Nature of Office Work and Workforce: Some of the network applications we have discussed would tend to alter markedly the nature of white-collar work and the knowledge and skills required of members of the workforce. That fact will give rise to issues involving reeducation and retraining, pay commensurate with responsibility, and displacement of labor by automation.

5) Impact on Productivity: Many people are expecting that applications such as electronic mail and office automation will significantly increase productivity, but there are as yet few if any definitive experiences with such applications or quantitative models of them that will convince skeptics. Impact-on-productivity may become a major issue in and of itself.

6) Educational Applications of Networks: Packet-switched and satellite networks, together with the great advances being made in computers, appear to open the door to revolutionary improvements in education, but much more than mere access and mere hardware will be required to achieve truly significant results. The issue that is arising is whether the society values education enough to support the long and difficult effort that will be required to develop effective computer- and network-based methods—or whether there will be another wave of premature exploitation followed by disappointment as there was in the computer-assisted-instruction "revolution" of the early 1960's.

7) Networks Versus Stand-Alone Systems: Why do we need time sharing when everyone can have his or her own microcomputer? What good is a network when one can have a whole library on a video disk? Those questions have answers, of course, in such applications as electronic message systems, distributed but cooperating "offices of the future," and computerized commerce, but the questions will nevertheless constitute a major issue. Microcomputers and inexpensive digital storage devices have significantly changed the network concept. Less than a decade ago, a computer network was something that provided access to a time-sharing system. Now it is a facility to support communication among spatially distributed people and computers and to supply people and computers with common information bases and supplementary storage and processing capabilities.

8) World Leadership: An important latent issue is implicit in the fact that different people have quite different perceptions of the importance of networking. A significant fraction of the people who have had experience as developers or intensive users of a packet-switching network believe they have been in on the beginning of a new era and that descendants of the ARPANET will constitute the nervous system of the world. On the other hand, most of the people who now determine the kind of national policy that earlier fostered the merchant marine, the railroads, the airlines, and the interstates seem not to be aware that any significant new potential exists or that there may be any reason to move rapidly to take advantage of it. And, of course, if it is meaningful at all to the man in the street, the term "information network" still suggests the telephone system, the radio, or a television network. In that situation, it is difficult to project as an issue the importance of networks to world economic leadership. We believe, nevertheless, that it is such an issue, and we hope that it will soon be recognized as such an issue.

9) Totalitarian Control: If almost all the telecommunications in an area were based on computerized networks controlled by an organization—say by a government—then, in the absence of effective safeguards, that organization could map the life space of every individual and record the business transactions of every company. The notion of telecommunications in the hands of a "big brother with computers" goes beyond the bounds of what is usually called "invasion of privacy" into the realm of totalitarian control. One can detect at least a trace of the "big brother" issue in the coldness of certain members of Congress toward, and the rejection by the OMB of, the plans (mentioned earlier) of the IRS to develop a computer- and network-based Tax Administration System. The mere possibility of subversion was enough to kill the system. It is of the utmost importance, of course, to develop truly effective safeguards against misuse of networks for purposes of social control. But such safeguards will be more difficult to devise than safeguards against ordinary invasion of privacy or against fraud and theft. Networks will have to be designed in such a way that representatives of diverse interests can satisfy themselves that there is no subversion and that the audit trails are not dossiers. And the arrangements will have to be dictator-proof. We think that that is a very great task and that it is being neglected.

V. Conclusions

Shakespeare could have been foreseeing the present situation in information networking when he said," . . . What's past is prologue; what to come, in yours and my discharge" [48]. Most of the applications that will shape the future of networking are now in the stage of conceptualization or in the stage of early development. But it seems possible that a "network of networks" will, even in this century, become the nervous system of the world and that its applications will significantly change the way we live and work. The degree to which the potentials of networking will be realized will depend upon how we resolve some of the issues that have been discussed.

The value of information networks will depend critically upon their connectivity and their ability to connect any one of many sources to any one or more of many destinations. High connectivity will be precluded if conditions force the development of many separate, independent, incompatible networks. One condition that would force such an incoherent development is the combination of 1) a need for security against loss of "electronic funds" and (other) proprietary

information and 2) the lack of a technology capable of providing security in an interconnected network or network of networks. That would lead to what we have called "fragmentation." Another such condition is based in a similar way on a combination of need for informational privacy and lack of the technology necessary to protect it except by isolating the privacy-sensitive data. The rapid and intensive development of computer and network security technology is vital to many network applications.

Many forces are fostering the development of networks to interconnect organization or branches of organizations and the development of applications to serve organizations, but there are few forces that foster networks to interconnect individuals or network applications to serve individuals. Perhaps the main hope for the provision of network services to individuals—especially network services to individuals at home—is that the Bell System will move (as obviously it would like to do) into the processing, storage, and information-commodity parts of the overall information business. But the telephone companies will be very slow to provide high-information-rate services because they have such large investments in narrow-band facilities. To get inexpensive wide-band channels into homes at an early date, we need a new departure in cable (or fiber optics) communication, taking off from cable television, or something truly revolutionary like a nation-wide network of aerostationary platforms: microwave platforms at 70 000 ft, supported by helium plus helicopter vanes, and relaying signals from housetop "dishes" a meter in diameter.

Examination of 30 actual and potential applications of networks suggests that the following network characteristics or capabilities are especially important: bidirectionality, freedom from undetected errors, efficiency despite "burstiness" in the transmission pattern, inherently low cost, high connectivity, high information transmission rate, security, privacy, authentication, and reliability. Mobility and broadcast capability turned out to be of the lowest priority in our analysis. Packet-switching and time-division-multiple-access networks, especially such networks with satellite relays, were suggested by the analysis to have the patterns of characteristics required to serve best the full range of applications. The analysis suggests an approach to the selection of the best network to serve any specified application or set of applications.

REFERENCES

[1] T. H. Myer and C. D. Mooers, *HERMES User's Guide*. Cambridge, MA: Bolt Beranek and Newman, Inc., June 1976.

[2] A. Vezza and M. S. Broos, "An electronic message system: Where does it fit?" in *Proc. IEEE Conf. Trends and Applications 1976: Computer Networks*, New York, Nov. 1976.

[3] A. Vezza, "A model for an electronic postal system," in *Telecommunications Policy Research Conf. Proc.*, Bruce M. Owen, Ed. Palo Alto, CA: Aspen Institute Program on Communications and Society, Apr. 1975, pp. 146-153.

[4] *SIGMA Message Service: Reference Manual*, Information Sciences Inst., Univ. Southern California, Marina Del Ray, CA, ISI/TM-78-11.

[5] J. B. Rule, "Value choices in electronic funds transfer policy," Office of Telecommunications Policy, Executive Office of the President, Washington, DC, Rep. Domestic Council Committee on the Right of Privacy, Oct. 1975.

[6] T. M. Dunn and J. H. Mirrissey, "Remote computing—An experimental system. Part 1: External specifications," in *Proc., Spring Joint Computer Conf.*, 1974.

[7] J. M. Keller, E. C. Strum, and G. H. Yang, "Remote computing—An experimental system. Part 2: Internal design," in *Proc., Spring Joint Computer Conf.*, 1974.

[8] J. E. Sammet, *Programming Languages: History and Fundamentals*. Englewood Cliffs, NJ: Prentice-Hall, 1969, pp. 226-229.

[9] T. Marill and D. Stern, "The datacomputer—A network data utility," in *AFIPS Conf. Proc.*, Montvale, NJ: AFIPS Press, May 1975.

[10] L. G. Roberts and B. D. Wessler, "Computer network development to achieve resource sharing," *AFIPS Conf. Proc.*, May 1970.

[11] L. G. Roberts, "ARPA network implications," ARPA Network Information Center Doc., 12 982, 1972.

[12] R. M. Metcalfe, "Packet communication," M.I.T. Lab. for Computer Science, MAC/TR-114, Dec. 1973.

[13] "SNDMSG," Tenex User's Guide, Bolt Beranek and Newman, Inc., Cambridge, MA, Oct. 1977.

[14] J. Vallee, "The forum project: Network conferencing and its future applications," *Comput. Networks (Netherlands)*, vol. 1, no. 1, pp. 39-52, June 1976.

[15] M. Turoff, "Party-line and discussion: Computerized conference systems," Office of Emergency Preparedness, Washington, DC, Rep., Apr. 1972.

[16] —— "Delphi conferencing. Computer-based conferencing with anonymity," in *Technological Forcasting and Social Change*. Washington, DC: National Resource Analysis Center, 1972, vol. 3, no. 2, pp. 159-204.

[17] M. Turoff and J. Scher, "Computerized conferencing and its impact on engineering management," in *Proc. 23rd Annu. Engineering Management Conf. on Effective Management of Engineering Resources*, Washington, DC, Oct. 1975.

[18] D. C. Engelbart, "Coordinated information services for a discipline- or mission-oriented community," in *Computer Communication Networks*, R. L. Grimsdale and F. F. Kuo, Eds. The Netherlands: Noordhoff International, 1975, pp. 89-99.

[19] D. C. Engelbart, R. W. Watson, and J. C. Norton, "The augmented knowledge workshop," in *AFIPS Conf. Proc.*, vol. 42, New York, June 1973.

[20] ——, *Advanced Intellect-Augmentation Techniques*. Menlo Park, CA: Stanford Research Institute, Feb. 1972.

[21] ——, "Human intellect augmentation techniques," Stanford Research Institute, Menlo Park, CA, Final Rep., 1969.

[22] R. Flesch, *The Art of Plain Talk*. New York: Harper & Brothers, 1946.

[23] J. O. Mason, Jr., "Management-information-systems—the auditor's role," *The Internal Auditor*, vol. 32, no. 5, pp. 40-48, Sept./Oct. 1975.

[24] S. J. Pokemper, "Management information-systems—A pragmatic survey," *Conf. Board Rec.*, vol. 10, no. 5, pp. 49-54, May 1973.

[25] R. G. Schroeder, "A survey of management science in university operations," *Manage. Sci.—Appl.*, vol. 19, no. 8, pp. 895-906, Apr. 1973.

[26] P. A. Strassman, "Managing the costs of information," *Harvard Bus. Rev.*, vol. 54, no. 5, pp. 133-142, Sept.-Oct. 1976.

[27] G. H. Barnes, R. M. Brown, M. Kato, D. J. Kuck, D. L. Slotnick, and R. A. Stoker, "The Illiac IV computer," *IEEE Trans. Comput.*, vol. C-17, pp. 746-757, 1968.

[28] W. J. Bouknight, S. A. Denenberg, D. E. McIntyre, J. M. Randall, A. H. Sameh, and D. L. Slotnick, "The Illiac IV system," *Proc. IEEE*, vol. 60, pp. 369-388, 1972.

[29] "Self-teaching design and analysis tools offered on timeshared basis," *Comput. Des.*, vol. 17, no. 2, pp. 18-20, Feb. 1978.

[30] I. E. Sutherland, *Computerized Commerce*. Santa Monica, CA: The Rand Paper Series, Rand Corp., Sept. 1975.

[31] National Research Council, Telecommunication for metropolitan areas: Near-term needs and opportunities. Washington, DC: Nat. Acad. Sci., 1977.

[32] L. L. Weed, "Technology is a link, not a barrier, for doctor and patient," *Modern Hospital*, vol. 114, pp. 80-83, Feb. 1970.

[33] ——, *Medical Records, Medical Education, and Patient Care. The Problem-Oriented record as a Basic Tool*. Chicago, IL: Year Book Medical Publishers, Inc., 1970.

[34] R. Davis, B. Buchanan, and E. Shortliffe, "Production rules as a representation for a knowledge-based consultation program," Stanford Univ. Artificial Intelligence Lab., Stanford, CA, Memo AIM-266, Oct. 1975.

[35] H. Silverman, "A digitalis therapy advisor," M.I.T. Lab. for Computer Science, MAC/TR-143, Jan. 1975.

[36] "GECOS: General Comprehensive Operating Supervisor," Honeywell Information Systems, Waltham, MA, Rep. DD19.

[37] A. Stathopoulos and H. F. Caley, "The Autodin II network," in *EASCON-77 Rec. Proc. IEEE Electronics and Aerospace Systems Convention* (Arlington, VA), Sept. 1977.

[38] Social Security Administration, Office of Advanced Systems, *Recommended Design Concept for the Future SSA Process*, Washington, DC, OAS Pub. No. 005, Apr. 1977.

[39] V. Slamenka, private communication, May 1978.

[40] CCITT, "Recommendation X.25: Interface between data terminal equipment (DTE) and data circuit-terminating equipment (DCE) for terminals operating in the packet mode on public data networks," in *Public Data Networks, Orange Book* (vol. VII.2). Geneva, Switzerland: Sixth Plenary Assembly, International Telecommunications Union, 1977, pp. 70-108.

[41] W. Diffie and M. Hellman, "New directions in cryptography," *IEEE Trans. Inform. Theory*, pp. 644-654, Nov. 1976.

[42] R. Rivest, A. Shamir, and L. Adleman, A method for obtaining digital signatures and public-key cryptosystems," M.I.T. Lab. for Computer Science, LCS/TM-82, Apr. 1977.
[43] *Federal Register*, vol. 40, no. 52, Mar. 17, 1975.
[44] *Federal Register*, vol. 40, no. 149, Aug. 1, 1975.
[45] W. Diffie and M. Hellman, "Exhaustive cryptanalysis of the NBS data encryption standard," *Computer*, vol. 10, no. 6, pp. 74-84, June 1977.
[46] J. H. Saltzer and M. D. Schroeder, "The protection of information in computer systems," *Proc. IEEE*, vol. 63, pp. 1278-1308, Sept. 1975.
[47] Office of Management and the Budget, OMB Circular A-108, Washington, DC, July 1975.
[48] William Shakespeare, *The Tempest*, act II, scene I.

A Tutorial on Protocols

LOUIS POUZIN AND HUBERT ZIMMERMANN

Invited Paper

Abstract—Protocols are common tools for controlling information transfer between computer systems. The concept of a protocol, which grew out of experimental computer networking, is now fundamental to system design. In this paper, basic protocol functions are explained and discussed. Then, the concept of a distributed system architecture is presented. It provides the framework for layers and protocols to operate across heterogeneous systems. The purpose and functions of each protocol layer such as, transmission, transport, virtual terminal, are described. Interactions between design and performance are discussed, and typical mechanisms are reviewed. CCITT and ISO relevant standards are summarized. Finally, the similarity between protocols and programming languages is emphasized as it points to the major impact brought about by protocols in system design.

I. Introduction

EVERYONE has had the opportunity to overhear such cryptic conversations exchanged over the radio by taxi drivers, policemen, and aircraft pilots. Although upon hearing these conversations at first do not mean much to the layman, these abbreviated languages carry well-defined meanings and obey well-defined rules. Speakers give their name, ask correspondents if they are listening, confirm reception, etc. This form of conversation differs drastically from a face-to-face chat. The communication channel is shared by many speakers. To save bandwidth and reduce interferences, messages are short and coded. External noise and other interferences are common occurences, hence, repetition and confirmation are normal practice. These rules are known as *protocols*.

The term *protocol* entered the computer jargon at the turn of the 70's, when the U.S. Defense Advanced Research Project Agency set out to build a network of geographically distribute heterogeneous computers [71]. Up to that time, communication between computer programs or processes was limited to processes which were located within the same machine. Interprocess communication was accomplished through the use of shared memory and special signals exchanged through the mediation of the operating system. This technique represented the analog of a face-to-face chat between processes. Interprocess communication between geographically distant systems would have left processes with the same kind of constraints that taxi drivers encounter. They would have to interact through a potentially hostile environment with limited bandwidth, delay and unreliable transmission. In addition, the processes in the different computer systems did not even speak the same native tongue, having been created by different manufacturers.

Computer veterans remember the sinuous evolution that led from binary programming to assembly code, to Fortran, Cobol, Algol, and other high-level languages. Originally viewed as a collection of tricks and hobbies, programming languages have developed into a major branch of computer science. The evolution of protocols has followed a strikingly similar path. Indeed, *Protocols are common tools designed for controlling information transfer between computer systems.* They are made up of sequences of messages with specific formats and meanings. These messages are equivalent to the instructions of a programming language, although protocol languages are still in an early stage of ad hoc development.

Manuscript received February 27, 1978; revised July 24, 1978.
The authors are with IRIA, 78150 Rocquencourt, France.

* got through M, N, & O. in the Encyclopedia.
read so much.
encouraged that curiosity.
sweete page @ age 16 —
the never discouraged the
curiosity element.